Video-R1: Reinforcing Video Reasoning in MLLMs

Kaituo Feng¹, Kaixiong Gong¹, Bohao Li², Zonghao Guo^{3*}, Yibing Wang⁴, Tianshuo Peng¹, Junfei Wu⁴, Xiaoying Zhang⁵, Benyou Wang², Xiangyu Yue^{1*}

¹CUHK MMLab, ²CUHK (SZ), ³Tsinghua University, ⁴UCAS, ⁵CUHK HCCL

 $\verb|https://github.com/tulerfeng/Video-R1||$

Abstract

Inspired by DeepSeek-R1's success in eliciting reasoning abilities through rulebased reinforcement learning (RL), we introduce Video-R1 as the first attempt to systematically explore the R1 paradigm for incentivizing video reasoning within multimodal large language models (MLLMs). However, directly applying RL training with the GRPO algorithm to video reasoning presents two primary challenges: (i) a lack of temporal modeling for video reasoning, and (ii) the scarcity of highquality video-reasoning data. To address these issues, we first propose the T-GRPO algorithm, which encourages models to utilize temporal information in videos for reasoning. Additionally, instead of relying solely on video data, we incorporate high-quality image-reasoning data into the training process. We have constructed two datasets: Video-R1-CoT-165k for SFT cold start and Video-R1-260k for RL training, both comprising image and video data. Experimental results demonstrate that Video-R1 achieves significant improvements on video reasoning benchmarks such as VideoMMMU and VSI-Bench, as well as on general video benchmarks including MVBench and TempCompass, etc. Notably, Video-R1-7B attains a 37.1% accuracy on video spatial reasoning benchmark VSI-bench, surpassing the commercial proprietary model GPT-4o. All code, models, and data are released in https://github.com/tulerfeng/Video-R1.

1 Introduction

Recent advancements in rule-based Reinforcement Learning (RL) [17] have significantly enhanced the reasoning capabilities of Large Language Models (LLMs) [16, 11]. In particular, DeepSeek-R1 [11] has demonstrated that carefully designed RL pipelines can lead to emergent and robust reasoning abilities with long chain-of-thoughts (CoT) in text-based domains. Motivated by this success, several recent efforts have explored extending RL training to Multimodal Large Language Models (MLLMs) [33, 14, 39]. Notable examples include Kimi k1.5 [33] and Skywork R1V [39], which apply RL to improve reasoning over image-text pairs. However, despite these early explorations, the domain of video reasoning in MLLMs remains underexplored.

To bridge this gap, we present Video-R1 model, as the first attempt to systematically investigate eliciting strong video reasoning based on the R1 paradigm . However, directly applying RL training with the Group Relative Policy Optimization (GRPO) algorithm [30] to video reasoning introduces two fundamental challenges: First, original GRPO lacks explicit reward signals for encouraging temporal reasoning in a video. Without explicit temporal awareness, the model may take shortcuts for reasoning, focusing on a single frame or snapshot rather than reasoning over time (see Figure 1 for example; Video-UTR [40] also identifies a similar issue). The underexploitation for temporal cues can cause the learned reasoning strategies to "shortcut" the process—relying on superficial visual patterns, rather than engaging in deeper and temporally grounded reasoning. This could ultimately hindering generalization to more complex or diverse video reasoning tasks.

^{*}Corresponding Authors

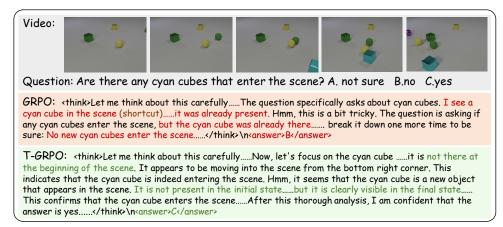


Figure 1: Reasoning paths of Video-R1 trained by GRPO and our proposed T-GRPO on test samples. Without explicit temporal modeling, models may learn sub-optimal video reasoning patterns by taking shortcuts, therefore failing to generalize well.

The second issue lies in the scarcity of high-quality video reasoning training data, especially samples that demand strong reasoning ability or involve long reasoning path. Most existing video training datasets mainly focus on simple recognition tasks, rather than reasoning. This scarcity makes it difficult to expose the model to diverse, challenging reasoning patterns during training, limiting the effectiveness of RL and hindering the emergence of robust reasoning behaviors.

To address these challenges, we propose two key solutions. First, we propose T-GRPO, an extension of the original GRPO algorithm [30] that explicitly encourages temporal reasoning. During training, the model is presented with both temporally ordered and randomly shuffled frame sequences, producing two groups of responses. A positive reward is assigned only when the proportion of correct answers from the ordered group exceeds that from the shuffled one. This strategy encourages the model to exploit temporal reasoning policy rather than relying on shortcuts derived from isolated frames.

Besides, to tackle the scarcity of high-quality video reasoning data, we strategically introduce image-based reasoning data as part of training data. We construct two datasets: Video-R1-CoT-165k for SFT cold start and Video-R1-260k for RL training. The image data serves as a valuable foundation for training general reasoning skills, while the curated video samples provide the temporal complexity needed for video understanding. This hybrid training setup not only alleviates the data bottleneck but also enables the model to transfer reasoning skills learned from static images to dynamic video contexts. Combined with T-GRPO, this approach equips Video-R1 with stronger, more generalizable video reasoning capabilities.

Our experiments show that Video-R1 achieves consistent and significant improvements across a suite of challenging video reasoning benchmarks, including VSI-Bench [38], VideoMMMU [13], MMVU [48], MVBench [20], TempCompass [27], and VideoMME [9]. Notably, Video-R1-7B attains 37.1% accuracy on VSI-Bench, a challenging video spatial reasoning benchmark, outperforming even proprietary models like GPT-40 [15]. These results suggest that with carefully designed algorithms and data pipelines, RL can indeed unlock complex temporal reasoning capabilities in MLLMs, similar to the breakthroughs seen in the text domain. Our contributions can be summaried as follows:

- We propose Video-R1, as the first attempt to systematically to explore developing video reasoning MLLMs based on the R1 paradigm. To support training, we construct two reasoning datasets: Video-R1-CoT-165k for SFT and Video-R1-260k for RL training, incorporating both image and video reasoning samples. We hope that Video-R1 will serve as a foundation for future research on video reasoning.
- To address the lack of temporal modeling in existing RL methods, we introduce **T-GRPO**, a novel training algorithm that encourages the model to utilize temporal information by contrasting reasoning performance over ordered and shuffled video frames.
- Extensive experiments on multiple video benchmarks, such as VideoMMMU, VSI-Bench, MVBench, etc, demonstrate the effectiveness of our approach. Notably, Video-R1-7B achieves 37.1% accuracy on VSI-Bench, outperforming the proprietary GPT-40 model.

2 Related Works

2.1 Multimodal Large Language Models for Video

Video understanding is an essential capability for Multimodal Large Language Models (MLLMs), enabling them to interpret and reason over dynamic visual content [31, 4, 37, 44, 46, 32, 2]. In recent years, a number of MLLMs have been developed specifically to advance progress in video understanding tasks. For example, LLaMA-VID [22] proposes a dual-token strategy (context and content tokens) to compress video input representations, enabling vision-language models to efficiently handle long videos while retaining essential visual information. VideoLLaMA2 [4] enhances video-language modeling by introducing spatial-temporal convolution for better dynamic understanding and an audio branch to integrate multimodal cues for richer video comprehension. LongVA [43] extends the context window of language backbones to process significantly longer video sequences without specialized video training, offering a language-centric solution to long-range temporal reasoning. VISA [37] introduces a knowledge-driven video object segmentation task that combines world knowledge with object tracking, addressing implicit, complex video queries through a segmentation-enabled multimodal LLM. These advancements highlight the potential of MLLMs in advancing video understanding. However, most prior works have primarily focused on video perception tasks. The development of MLLMs with strong video reasoning capabilities remains largely unexplored.

2.2 Large Language Model Reasoning

The reasoning abilities of Large Language Models (LLMs) have been a focal point of recent research, aiming to enhance their capacity to perform complex, multi-step problem-solving tasks [34, 47, 42, 49, 23, 8, 35, 41, 24]. Unlike earlier approaches that rely on dense, step-level supervision or learned reward models to supervise reasoning paths [10, 21], DeepSeek-R1 initiates a new wave of interest in rule-based reinforcement learning, demonstrating that even coarse, outcome-only rewards can effectively elicit strong reasoning behavior [11]. Its success shows that with a carefully designed reward structure and policy optimization strategy, models could learn to generate long CoT without requiring intermediate supervision. Following this paradigm, several recent efforts have attempted to reproduce R1's success [14, 33, 36, 3, 28, 18, 6, 45]. For example, Open Reasoner Zero [12] and Kimi k1.5 [33] explore similar rule-based RL pipelines to enhance reasoning in the text and image domains, respectively. However, despite encouraging progress, few prior work has explored how to extend this approach to the video domain. Bridging this gap remains an open challenge and a promising direction for expanding the boundaries of reasoning models.

3 Methods

3.1 Dataset Construction

High-quality training data plays a crucial role in reinforcing video reasoning capabilities in MLLMs. In this section, we will introduce how we curate Video-R1-260k for RL training and Video-R1-CoT-165k for SFT cold start.

Data Collection and Curation. To overcome the scarcity of high-quality video reasoning training data, we strategically introduce image-based reasoning data as part of training data. The image-based data serves primarily to teach the model a broad range of reasoning skills, covering various difficulty levels and domains such as math, spatial logic, expert-level knowledge, etc. These samples help the model develop generalized reasoning abilities in static contexts. In contrast, the video-based data is primarily used to train the model's ability to perform temporal reasoning—including understanding event progression, capturing frame-to-frame dependencies, and drawing inferences based on motion and causal dynamics over time.

We collect data from a variety of public datasets and carefully sample and balance the proportion of each subset. The final composition of the Video-R1-260k dataset is illustrated in Figure 2. The distribution of Video-R1-260k dataset can be roughly categorized as follows:

• General (Video, 116k): A diverse set of open-domain video data, covering everyday scenarios and designed to build temporal comprehension and reasoning abilities.

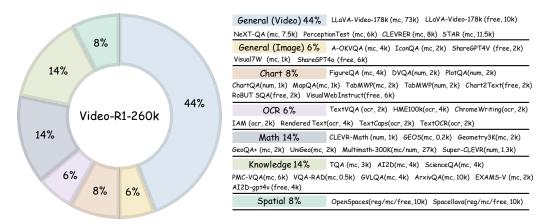


Figure 2: The data distribution of our Video-R1-260k dataset.

- General (Image, 15k): A general-purpose image question-answering data, used to provide basic visual understanding.
- Chart (Image, 21k): Visual reasoning over charts, line graphs and scientific figures, focusing on data interpretation and quantitative logic.
- OCR (Image, 16k): Facilitate reasoning tasks that require recognizing and interpreting embedded textual content such as signs, forms, or documents.
- Math (Image, 37k): Image-based math reasoning questions, including formulas, geometry diagrams, and multi-step symbolic reasoning.
- **Knowledge** (**Image**, **37k**): Visual commonsense and multi-discipline reasoning tasks, testing the model's ability to integrate world knowledge with visual cues.
- Spatial (Image, 20k): Tasks that require understand spatial information for reasoning.

CoT Annotation. To facilitate an effective SFT cold start, we leverage Qwen2.5-VL-72B-Instruct [1] to generate CoT rationales for the samples in Video-R1-260k. The prompt template for CoT generation is provided in Appendix C.1. After applying basic rule-based filtering to remove low-quality or inconsistent outputs, we obtain a high-quality CoT dataset, Video-R1-CoT-165k, which is used for the cold-start SFT stage.

Data Type and Rule-based Reward Design. Since our reinforcement learning framework follows the rule-based reward paradigm of DeepSeek-R1 [11], it is crucial to ensure that the reward signals are both reliable and precise. To this end, the majority of our training data is designed around tasks with clearly verifiable outputs, such as multiple-choice and numerical answer formats. This allows for accurate reward computation using simple rules, ensuring stable and effective RL training.

However, to increase the model's flexibility and its ability to generalize across diverse tasks and formats, we also incorporate a smaller portion of other data types. These include free-form generation, OCR tasks, and regression problems, which are essential for adapting to real-world applications and broader datasets.

The data types and corresponding correctness reward functions are summarized as follows:

- **Multiple Choice:** The reward is assigned based on whether the predicted answer matches the ground-truth answer.
- Numerical QA: A binary reward is given depending on whether the predicted number exactly matches the reference value.
- OCR: We compute the reward using the Word Error Rate (WER), measuring the edit distance between the predicted and reference text.
- Free-form QA: The reward is calculated as the average of ROUGE-1, ROUGE-2, and ROUGE-L scores between the model's output and the ground-truth answer.
- **Regression:** The closer the predicted value is to the ground truth, the higher the reward, calculated as one minus their relative error.

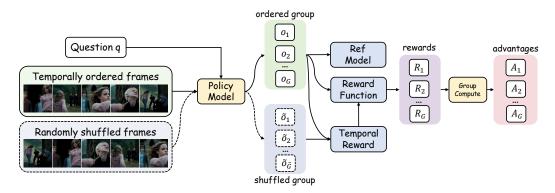


Figure 3: An illustration of our proposed T-GRPO algorithm.

3.2 Temporal Group Relative Policy Optimization (T-GRPO)

While GRPO [30] has proven effective in text-based reasoning, it lacks explicit reward signals for temporal reasoning—making it insufficient for training MLLMs to reason over videos. To address this, we propose Temporal Group Relative Policy Optimization (T-GRPO), which introduces a contrastive reward mechanism that explicitly encourages temporal reasoning, as illustrated in Figure 3.

The core idea behind T-GRPO is to compare the model's performance on the same video question when frames are provided in two different orders: (1) the *temporally ordered* sequence, and (2) a *randomly shuffled* version. For each input question, we generate two groups of responses $\{o_i\}_{i=1}^G$ and $\{\tilde{o}_i\}_{i=1}^{\tilde{G}}$ using the ordered and shuffled frame inputs, respectively.

Let p and \tilde{p} denote the proportion of correct answers in each group. We then define a temporal reward r_t as:

$$r_t = \begin{cases} \alpha, & \text{if } p \ge \tilde{p} \\ 0, & \text{otherwise} \end{cases} \tag{1}$$

where α is a hyperparameter controlling the magnitude of the temporal reward. Here we set $\alpha = 0.3$.

This contrastive design encourages the model to perform better when the video is presented in correct temporal order than when it is shuffled. The model is only granted this positive reward if its current reasoning strategy for a given question demonstrates a reliance on temporal information. For tasks with continuous rewards (e.g., free-form answers), a threshold (e.g., 0.5) can be used to determine whether a response is considered correct.

Importantly, r_t is only applied to correct responses to ensure meaningful positive advantages. Applying it to all responses would dilute the reward signal and hinder effective learning. In other words, when the model's reasoning policy successfully relies on temporal patterns, correct responses are reinforced with a higher reward, while incorrect ones remain unaffected.

Formally, the temporal-augmented reward is defined as:

$$R_i = \begin{cases} r_i + r_t, & \text{if } o_i \text{ is correct} \\ r_i, & \text{otherwise} \end{cases}$$
 (2)

where r_i is the reward for response i, containing both the correctness reward and the format reward, following [11]. R_i is the final reward used for calculating advantages. This reward shaping ensures that when the model answers correctly under a temporal setting but fails to outperform the shuffled baseline, it receives no additional reward—pushing the optimization toward adopting a more temporally aware reasoning policy. The temporal reward r_t could also be added to the advantages directly.

Then, the advantage A_i is computed over the rewards within each group:

$$A_i = \frac{R_i - \text{mean}(\{R_j\})}{\text{std}(\{R_i\})}$$
(3)

Following DeepSeek R1 [11], the final policy update is as follows:

$$\mathcal{J}_{\text{T-GRPO}}(\theta) = \mathbb{E}_{q,\{o_i\}} \left[\frac{1}{G} \sum_{i=1}^{G} \left(\min \left(\frac{\pi_{\theta}(o_i|q)}{\pi_{\theta_{\text{old}}}(o_i|q)} A_i, \quad \text{clip} \left(\frac{\pi_{\theta}(o_i|q)}{\pi_{\theta_{\text{old}}}(o_i|q)}, 1 - \epsilon, 1 + \epsilon \right) A_i \right) - \beta \, \mathbb{D}_{\text{KL}}(\pi_{\theta} || \pi_{\text{ref}}) \right) \right]$$

$$(4)$$

By explicitly comparing the model's performance under ordered and shuffled inputs, T-GRPO introduces a contrastive training signal that drives the model to prefer reasoning strategies that leverage temporal patterns. It is worth noting that T-GRPO is only employed for video-based inputs in the training process of Video-R1.

3.3 Training Strategies

We adopt Qwen2.5-VL-7B-Instruct [1] as the base MLLMs for training. Similar to DeepSeek-R1 [11], the training process is conducted in two stages: SFT cold start followed by RL training. For these two stages, we both adopt image-video mixed training strategy.

In the first stage, we perform SFT on the Video-R1-CoT-165k dataset. This step serves as a cold-start initialization, equipping the model with basic reasoning capabilities across a variety of modalities. The resulting model is denoted as Qwen2.5-VL-7B-SFT.

In the second stage, we further train the Qwen2.5-VL-7B-SFT model on the broader Video-R1-260k dataset using our proposed T-GRPO algorithm. This reinforcement learning phase is designed to guide the model beyond the rigid, pattern-matching behavior induced by SFT, encouraging it to freely explore better reasoning strategies. The resulting model is denoted as Video-R1-7B.

To further enhance the quality of reasoning, we introduce a length-based reward to regulate the length of the model's output. Specifically, this mechanism aims to strike a balance between encouraging deeper reasoning and preventing overthinking. For each reasoning path o_i , if the predicted answer is correct and the response length falls within a predefined interval $[l_{\min}, l_{\max}]$, the model receives an additional reward $r_l = \omega$. Formally:

$$R_i = \begin{cases} R_i + \omega, & \text{if } o_i \text{ is correct and } l_{\min} \leq \text{len}(o_i) \leq l_{\max} \\ R_i, & \text{otherwise} \end{cases}$$
 (5)

This reward encourages the model to think deeply without overthinking. In this paper, we set $\omega=0.2$, $l_{min}=320$ and $l_{max}=512$. The analysis of this reward's effect is provided in Appendix A.2.

3.4 Aha Moment in Video Reasoning

We find that Video-R1 is capable of performing complex reasoning with self-reflective behaviors, often referred to as "aha moments" [11]. Specifically, the model occasionally revisits its interpretation of the video or reflects on previous outputs, particularly when confronted with ambiguous temporal cues or multi-step inference tasks. These behaviors indicate that the model is not merely executing memorized patterns but is actively engaging in internal feedback loops by re-examining evidence. An example is demonstrated in Figure 4. More examples can be found in Appendix B.

4 Experiments

4.1 Setup

Benchmarks. We evaluate our model on six video benchmarks: VSI-Bench [38], VideoMMMU [13], MMVU [48], MVBench [20], TempCompass [27], and VideoMME [9]. Among them, the first three are video reasoning benchmarks, which focus primarily on assessing the model's reasoning capabilities in video understanding. The latter three are general-purpose video understanding benchmarks, which include a mixture of perception and reasoning tasks. For MMVU, we evaluate on its multiple-choice question set for stability and consistency. For all evaluations, we follow the decoding configuration used in the official Qwen2.5-VL demo, with top p = 0.001 and temperature = 0.01.

Video: Question: You are a robot beginning at the blue built-in bookshelf facing the blue built-in bookshelf. You want to navigate to the bathroom sink . You will perform the following actions (Note: for each [please fill in], choose either 'turn back,' 'turn left,' or 'turn right.'): 1. [please fill in] 2. Go forward until the door of the bathroom. 3. [please fill in] 4. Go forward until the sink. You have reached the final destination. A. Turn Back, Turn Right B. Turn Left, Turn Left C. Turn Right, Turn Left D. Turn Back, Turn Left Qwen2.5-VL-7B-Instruct: Let me think through this step by step. I start at the blue built-in bookshelf. To get to the bathroom sink, I need to move in a specific direction. The first action is to go forward, which means I need to move away from the bookshelf. Since I'm facing the bookshelf, I should turn left to move forward. After that, I need to go forward until I reach the bathroom door. Once I'm at the door, I should turn left to go into the bathroom and then turn left again to reach the sink. This sequence of actions aligns with option B. <answer>B</ans Video-R1-7B: think-Let me think about this carefully. The goal is to navigate from the blue built-in bookshelf to the bathroom sink. The first step is to determine the direction from the bookshelf to the bathroom 1.The blue built-in bookshelf is on the left side of the room. The bathroom is on the right side of the room. So, I need to turn right from the bookshelf to head towards the bathroom. 2. After turning right, I need to go forward until I reach the door of the bathroom. This means I need to continue moving forward in the direction I just turned 3. Once I reach the bathroom door, I need to turn left to enter the bathroom. This is because the sink is typically located on the left side of the bathroom. So, the sequence of actions should be:1. Turn right from the bookshelf. 2. Go forward until the bathroom door. 3. Turn left to enter the bathroom. 4. Go forward until the sink. (Perception and Reasoning) Hmm this seems to alian with the sequence of actions in option C. Let's break it down one more time to verify Option A: Turn Back, Turn Right - This would mean going back to the bookshelf and then turning right, which doesn't make sense for reaching the bathroom Option B: Turn Left, Turn Left - This would mean turning left from the bookshelf and then turning left again, which also doesn't make sense Option D: Turn Back, Turn Left - This would mean going back to the bookshelf and then turning left, which is not the correct path. (Reflection) Oh, I see. The correct sequence is indeed in option C. Let's confirm one last time. The correct sequence is to turn right, go forward to the bathroom door, turn left, and go forward to the sink. This matches option C. Final answer: C Let's verify one last time. Yes, C is the correct answer.</think> <answer>C</answer>

Figure 4: An example of Video-R1-7B's reasoning output on VSI-Bench.

Training Details. We train our model using up to 8 NVIDIA A100 (80GB) GPUs. For efficiency considerations, we limit the maximum number of video frames to 16 during training. Each frame is processed at a max resolution of $128 \times 28 \times 28$ pixels. During inference, we increase the frame resolution to $256 \times 28 \times 28$ pixels and frames to $16{\sim}64$ to enhance performance. The ordered group size G is set to 8 and the shuffled group size G is set to half of that for efficiency. More details are provided in Appendix D.

We first perform SFT on Video-R1-CoT-165k for one epoch to obtain the Qwen2.5-VL-7B-SFT model. This is followed by RL training on Video-R1-260k to produce the final Video-R1 model. Due to current computational resource limitations, we train the model for only 1k RL steps. Surprisingly, even within this limited training budget, the model exhibits significant improvements in video reasoning performance, indicating the strong effectiveness of both our data design and algorithm. Further experiments exploring the impact of scaling up RL training are presented in Appendix A.1.

4.2 Main Results

As shown in Table 1. Our experimental results across six benchmarks validate the effectiveness of Video-R1 in video reasoning and general video understanding tasks. The key findings are as follows.

Superior Performance of Video-R1. Video-R1 significantly outperforms previous models across most benchmarks, with particularly strong gains on video reasoning tasks such as VSI-Bench, VideoMMMU, and MMVU. This highlights the necessity of explicit reasoning capability in solving video tasks, and confirms the effectiveness of reinforcement learning for video tasks.

RL Works Better Than SFT. We observe that the SFT model Qwen2.5-VL-7B-SFT does not consistently improve performance. In some cases (e.g., VideoMME), performance even slightly drops after SFT, likely due to overfitting or limited generalization in unseen scenarios [5]. In contrast, after only 1k steps of reinforcement learning, Video-R1 achieves significant performance boosts. This clearly demonstrates the importance of RL in unlocking generalizable video reasoning capability.

Table 1: Performance of different models on benchmarks.

Models	Frames	Video Reasoning Benchmark			Video General Benchmark			
3.3340.0	77411100	VSI-Bench	VideoMMMU	MMVU (mc)	MVBench	TempCompass	VideoMME (wo sub)	
GPT-4o [15]	-	34.0	61.2	75.4	-	-	71.9	
LLaMA-VID [22]	-	-	-	-	41.9	45.6	-	
VideoLLaMA2 [4]	-	-	-	44.8	54.6	-	47.9	
LongVA-7B [43]	-	29.2	23.9	-	-	56.9	52.6	
VILA-1.5-8B [25]	-	28.9	20.8	-	-	58.8	-	
VILA-1.5-40B [25]	-	31.2	34.0	-	-	-	60.1	
Video-UTR-7B [40]	-	-	-	-	58.8	59.7	52.6	
LLaVA-OneVision-7B [19]	-	32.4	33.8	49.2	56.7	-	58.2	
Kangaroo-8B [26]	-	-	-	-	61.1	62.5	56.0	
Qwen2.5-VL-7B (CoT)	16	27.7	47.8	59.2	57.4	72.2	53.1	
Qwen2.5-VL-7B-SFT	16	31.8	47.4	61.3	59.4	69.2	52.8	
Qwen2.5-VL-7B (CoT)	32	30.1	48.1	60.0	59.0	72.6	56.6	
Qwen2.5-VL-7B-SFT	32	33.3	49.4	63.5	60.5	69.9	55.4	
Qwen2.5-VL-7B (CoT)	64	31.4	50.4	60.0	59.2	72.9	59.6	
Qwen2.5-VL-7B-SFT	64	34.8	49.4	61.6	60.6	70.0	58.8	
Video-R1-7B	16	34.6	49.8	64.2	62.7	72.6	57.4	
Video-R1-7B	32	35.8	52.3	63.8	63.9	73.2	59.3	
Video-R1-7B	64	37.1	52.4	63.8	64.8	73.2	61.4	

Table 2: Ablation Study.

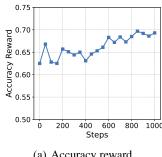
Models	Frames	Video Reasoning Benchmark			Video General Benchmark		
			VideoMMMU	MMVU (mc)	MVBench	TempCompass	VideoMME (wo sub)
Video-R1-7B-wo-image	16	32.3	45.8	60.6	60.9	69.8	53.8
Video-R1-7B-wo-temporal	16	32.7	48.3	62.1	61.1	71.3	54.5
Video-R1-7B-zero	16	31.8	49.5	63.8	60.4	70.9	53.8
Video-R1-7B	16	34.6	49.8	64.2	62.7	72.6	57.4

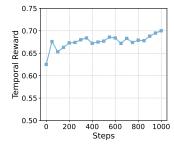
More Frames Lead to Better Reasoning. When increasing the input frame number, we observe performance improvements on almost all benchmarks. This indicates that richer context and temporal information contribute positively to the reasoning performance. Therefore, developing models capable of understanding and reasoning over longer video inputs is a promising direction for future research.

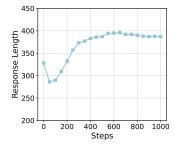
4.3 Ablation Study

In this section, we conduct an ablation study by designing three variants of our model: Video-R1-7B-wo-image, which removes all image-based data during training and relies solely on video data; Video-R1-7B-wo-temporal, which replaces our proposed T-GRPO algorithm with the original GRPO method; Video-R1-zero, which skips the SFT cold start and directly conducts RL training.

As shown in Table 2, both ablated models perform worse than the full Video-R1-7B across all benchmarks. In particular, removing image data leads to a noticeable drop in performance on both video reasoning and general benchmarks, indicating that image-based samples play a crucial role in bootstrapping general reasoning ability. Similarly, without temporal-aware training via T-GRPO, the model struggles to fully leverage temporal cues, resulting in weaker performance on benchmarks. Additionally, skipping the SFT cold start leads to worse performance, underscoring the value of our Video-R1-CoT-165k dataset and the importance of initializing with SFT before RL training. These ablations validate the effectiveness of our proposed methods.







(a) Accuracy reward

(b) Temporal reward r_t (scaled)

(c) Response length

Figure 5: RL training curves.

Training Curves

Figure 5 illustrates the RL training dynamics of Video-R1. As shown in Figure 5(a), the accuracy reward exhibits a generally upward trend, indicating that the model continuously improves its ability to produce correct answers under reinforcement learning. In Figure 5(b), the temporal reward r_t (scaled to $0\sim1$ for visibility) also demonstrates a steady increase. This suggests that the model is progressively adopting more temporally grounded reasoning strategies with T-GRPO. Interestingly, the response length in Figure 5(c) drops initially during RL training, then rises and stabilizes. We guess that this may reflect a learning transition: the model first discards its sub-optimal SFT reasoning style, and eventually settles on a new reasoning policy.

Effect of Temporal Reward Analysis

To further assess the impact of T-GRPO, we measure the percentage of responses that incorporate temporal reasoning for questions requiring it. Specifically, we use Qwen2.5-VL-72B to identify all temporally grounded questions across six benchmarks, and then use Qwen2.5-VL-72B to evaluate whether the model responses demonstrate temporal reasoning. The evaluation prompt can be

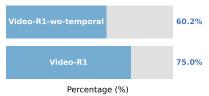


Figure 6: Percentage of temporal reasoning responses.

found in Appendix C.2. As shown in Figure 6, Video-R1 trained by T-GRPO incorporates temporal reasoning in 75.0% of responses, compared to 60.2% for Video-R1-wo-temporal, which lacks temporal modeling. This clear gap demonstrates the effectiveness of T-GRPO in encouraging the model to leverage temporal information during reasoning.

We also conduct a sensitivity analysis on the magnitude of the temporal reward, controlled by the hyperparameter α , as shown in Figure 7. The average performance is calculated as the mean performance across six benchmarks. We observe a slight drop in performance at $\alpha=0.1$ and $\alpha = 0.4$, while $\alpha = 0.2$ and $\alpha = 0.3$ yield similar and favorable results. This indicates that the model is relatively insensitive to the choice of α within a reasonable range.

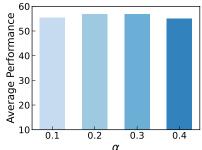


Figure 7: Sensitivity analysis.

Conclusions

In this work, we present Video-R1, as an attempt to investigate the R1 paradigm to enhances video reasoning in multimodal large language models. Motivated by DeepSeek-R1, we extend rule-based reinforcement learning to video by introducing T-GRPO, a temporal-aware extension of GRPO that explicitly encourages temporal reasoning in video. To support training, we curate two datasets, Video-R1-CoT-165k for SFT cold start and Video-R1-260k for RL. Experimental results across six benchmarks validate the effectiveness of our approach. We hope this work provides a foundation for further research in video reasoning with MLLMs.

6 Acknowledgment

This work is partially supported by the National Natural Science Foundation of China (Grant No. 62306261), and The Shun Hing Institute of Advanced Engineering (SHIAE) Grant (No. 8115074). This study was supported in part by the Centre for Perceptual and Interactive Intelligence, a CUHK-led InnoCentre under the InnoHK initiative of the Innovation and Technology Commission of the Hong Kong Special Administrative Region Government.

References

- [1] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025.
- [2] Joya Chen, Zhaoyang Lv, Shiwei Wu, Kevin Qinghong Lin, Chenan Song, Difei Gao, Jia-Wei Liu, Ziteng Gao, Dongxing Mao, and Mike Zheng Shou. Videollm-online: Online video large language model for streaming video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18407–18418, 2024.
- [3] Qiguang Chen, Libo Qin, Jinhao Liu, Dengyun Peng, Jiannan Guan, Peng Wang, Mengkang Hu, Yuhang Zhou, Te Gao, and Wanxiang Che. Towards reasoning era: A survey of long chain-of-thought for reasoning large language models. *arXiv preprint arXiv:2503.09567*, 2025.
- [4] Zesen Cheng, Sicong Leng, Hang Zhang, Yifei Xin, Xin Li, Guanzheng Chen, Yongxin Zhu, Wenqi Zhang, Ziyang Luo, Deli Zhao, et al. Videollama 2: Advancing spatial-temporal modeling and audio understanding in video-llms. *arXiv preprint arXiv:2406.07476*, 2024.
- [5] Tianzhe Chu, Yuexiang Zhai, Jihan Yang, Shengbang Tong, Saining Xie, Dale Schuurmans, Quoc V Le, Sergey Levine, and Yi Ma. Sft memorizes, rl generalizes: A comparative study of foundation model post-training. *arXiv preprint arXiv:2501.17161*, 2025.
- [6] Kaixuan Fan, Kaituo Feng, Haoming Lyu, Dongzhan Zhou, and Xiangyu Yue. Sophiavl-r1: Reinforcing mllms reasoning with thinking reward. *arXiv preprint arXiv:2505.17018*, 2025.
- [7] Jiazhan Feng, Shijue Huang, Xingwei Qu, Ge Zhang, Yujia Qin, Baoquan Zhong, Chengquan Jiang, Jinxin Chi, and Wanjun Zhong. Retool: Reinforcement learning for strategic tool use in llms. *arXiv preprint arXiv:2504.11536*, 2025.
- [8] Kaituo Feng, Changsheng Li, Xiaolu Zhang, Jun Zhou, Ye Yuan, and Guoren Wang. Keypoint-based progressive chain-of-thought distillation for llms. *arXiv preprint arXiv:2405.16064*, 2024.
- [9] Chaoyou Fu, Yuhan Dai, Yongdong Luo, Lei Li, Shuhuai Ren, Renrui Zhang, Zihan Wang, Chenyu Zhou, Yunhang Shen, Mengdan Zhang, et al. Video-mme: The first-ever comprehensive evaluation benchmark of multi-modal llms in video analysis. *arXiv preprint arXiv:2405.21075*, 2024.
- [10] Jiaxuan Gao, Shusheng Xu, Wenjie Ye, Weilin Liu, Chuyi He, Wei Fu, Zhiyu Mei, Guangju Wang, and Yi Wu. On designing effective rl reward at training time for llm reasoning. *arXiv* preprint arXiv:2410.15115, 2024.
- [11] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.
- [12] Jingcheng Hu, Yinmin Zhang, Qi Han, Daxin Jiang, and Heung-Yeung Shum Xiangyu Zhang. Open-reasoner-zero: An open source approach to scaling reinforcement learning on the base model, 2025.
- [13] Kairui Hu, Penghao Wu, Fanyi Pu, Wang Xiao, Yuanhan Zhang, Xiang Yue, Bo Li, and Ziwei Liu. Video-mmmu: Evaluating knowledge acquisition from multi-discipline professional videos. *arXiv preprint arXiv:2501.13826*, 2025.

- [14] Wenxuan Huang, Bohan Jia, Zijie Zhai, Shaosheng Cao, Zheyu Ye, Fei Zhao, Zhe Xu, Yao Hu, and Shaohui Lin. Vision-r1: Incentivizing reasoning capability in multimodal large language models. *arXiv preprint arXiv:2503.06749*, 2025.
- [15] Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. arXiv preprint arXiv:2410.21276, 2024.
- [16] Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. Openai o1 system card. arXiv preprint arXiv:2412.16720, 2024.
- [17] Leslie Pack Kaelbling, Michael L Littman, and Andrew W Moore. Reinforcement learning: A survey. *Journal of artificial intelligence research*, 4:237–285, 1996.
- [18] Yuxiang Lai, Jike Zhong, Ming Li, Shitian Zhao, and Xiaofeng Yang. Med-r1: Reinforcement learning for generalizable medical reasoning in vision-language models. *arXiv preprint arXiv:2503.13939*, 2025.
- [19] Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Peiyuan Zhang, Yanwei Li, Ziwei Liu, et al. Llava-onevision: Easy visual task transfer. *arXiv preprint arXiv:2408.03326*, 2024.
- [20] Kunchang Li, Yali Wang, Yinan He, Yizhuo Li, Yi Wang, Yi Liu, Zun Wang, Jilan Xu, Guo Chen, Ping Luo, et al. Mvbench: A comprehensive multi-modal video understanding benchmark. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22195–22206, 2024.
- [21] Wendi Li and Yixuan Li. Process reward model with q-value rankings. arXiv preprint arXiv:2410.11287, 2024.
- [22] Yanwei Li, Chengyao Wang, and Jiaya Jia. Llama-vid: An image is worth 2 tokens in large language models. In *European Conference on Computer Vision*, pages 323–340. Springer, 2024.
- [23] Zhong-Zhi Li, Duzhen Zhang, Ming-Liang Zhang, Jiaxin Zhang, Zengyan Liu, Yuxuan Yao, Haotian Xu, Junhao Zheng, Pei-Jie Wang, Xiuyi Chen, et al. From system 1 to system 2: A survey of reasoning large language models. *arXiv preprint arXiv:2502.17419*, 2025.
- [24] Zongzhao Li, Zongyang Ma, Mingze Li, Songyou Li, Yu Rong, Tingyang Xu, Ziqi Zhang, Deli Zhao, and Wenbing Huang. Star-r1: Spatial transformation reasoning by reinforcing multimodal llms. *arXiv preprint arXiv:2505.15804*, 2025.
- [25] Ji Lin, Hongxu Yin, Wei Ping, Pavlo Molchanov, Mohammad Shoeybi, and Song Han. Vila: On pre-training for visual language models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 26689–26699, 2024.
- [26] Jiajun Liu, Yibing Wang, Hanghang Ma, Xiaoping Wu, Xiaoqi Ma, Xiaoming Wei, Jianbin Jiao, Enhua Wu, and Jie Hu. Kangaroo: A powerful video-language model supporting long-context video input. *arXiv preprint arXiv:2408.15542*, 2024.
- [27] Yuanxin Liu, Shicheng Li, Yi Liu, Yuxiang Wang, Shuhuai Ren, Lei Li, Sishuo Chen, Xu Sun, and Lu Hou. Tempcompass: Do video llms really understand videos? *arXiv preprint arXiv:2403.00476*, 2024.
- [28] Zhaowei Liu, Xin Guo, Fangqi Lou, Lingfeng Zeng, Jinyi Niu, Zixuan Wang, Jiajie Xu, Weige Cai, Ziwei Yang, Xueqian Zhao, et al. Fin-r1: A large language model for financial reasoning through reinforcement learning. *arXiv preprint arXiv:2503.16252*, 2025.
- [29] Zhengxi Lu, Yuxiang Chai, Yaxuan Guo, Xi Yin, Liang Liu, Hao Wang, Guanjing Xiong, and Hongsheng Li. Ui-r1: Enhancing action prediction of gui agents by reinforcement learning. *arXiv* preprint arXiv:2503.21620, 2025.
- [30] Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Y Wu, et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.

- [31] Fangxun Shu, Lei Zhang, Hao Jiang, and Cihang Xie. Audio-visual llm for video understanding. *arXiv preprint arXiv:2312.06720*, 2023.
- [32] Yunlong Tang, Jing Bi, Siting Xu, Luchuan Song, Susan Liang, Teng Wang, Daoan Zhang, Jie An, Jingyang Lin, Rongyi Zhu, et al. Video understanding with large language models: A survey. *IEEE Transactions on Circuits and Systems for Video Technology*, 2025.
- [33] Kimi Team, Angang Du, Bofei Gao, Bowei Xing, Changjiu Jiang, Cheng Chen, Cheng Li, Chenjun Xiao, Chenzhuang Du, Chonghua Liao, et al. Kimi k1. 5: Scaling reinforcement learning with llms. *arXiv preprint arXiv:2501.12599*, 2025.
- [34] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022.
- [35] Junfei Wu, Jian Guan, Kaituo Feng, Qiang Liu, Shu Wu, Liang Wang, Wei Wu, and Tieniu Tan. Reinforcing spatial reasoning in vision-language models with interwoven thinking and visual drawing. *arXiv preprint arXiv:2506.09965*, 2025.
- [36] Tian Xie, Zitian Gao, Qingnan Ren, Haoming Luo, Yuqian Hong, Bryan Dai, Joey Zhou, Kai Qiu, Zhirong Wu, and Chong Luo. Logic-rl: Unleashing llm reasoning with rule-based reinforcement learning. *arXiv* preprint arXiv:2502.14768, 2025.
- [37] Cilin Yan, Haochen Wang, Shilin Yan, Xiaolong Jiang, Yao Hu, Guoliang Kang, Weidi Xie, and Efstratios Gavves. Visa: Reasoning video object segmentation via large language models. In *European Conference on Computer Vision*, pages 98–115. Springer, 2024.
- [38] Jihan Yang, Shusheng Yang, Anjali W Gupta, Rilyn Han, Li Fei-Fei, and Saining Xie. Thinking in space: How multimodal large language models see, remember, and recall spaces. *arXiv* preprint arXiv:2412.14171, 2024.
- [39] Chris Yi Peng et al. Skywork r1v: Pioneering multimodal reasoning with chain-of-thought, 2025.
- [40] En Yu, Kangheng Lin, Liang Zhao, Yana Wei, Zining Zhu, Haoran Wei, Jianjian Sun, Zheng Ge, Xiangyu Zhang, Jingyu Wang, et al. Unhackable temporal rewarding for scalable video mllms. *arXiv preprint arXiv:2502.12081*, 2025.
- [41] Jiakang Yuan, Tianshuo Peng, Yilei Jiang, Yiting Lu, Renrui Zhang, Kaituo Feng, Chaoyou Fu, Tao Chen, Lei Bai, Bo Zhang, et al. Mme-reasoning: A comprehensive benchmark for logical reasoning in mllms. *arXiv preprint arXiv:2505.21327*, 2025.
- [42] Lifan Yuan, Ganqu Cui, Hanbin Wang, Ning Ding, Xingyao Wang, Jia Deng, Boji Shan, Huimin Chen, Ruobing Xie, Yankai Lin, et al. Advancing llm reasoning generalists with preference trees. *arXiv preprint arXiv:2404.02078*, 2024.
- [43] Peiyuan Zhang, Kaichen Zhang, Bo Li, Guangtao Zeng, Jingkang Yang, Yuanhan Zhang, Ziyue Wang, Haoran Tan, Chunyuan Li, and Ziwei Liu. Long context transfer from language to vision. arXiv preprint arXiv:2406.16852, 2024.
- [44] Xiaoying Zhang, Da Peng, Yipeng Zhang, Zonghao Guo, Chengyue Wu, Chi Chen, Wei Ke, Helen Meng, and Maosong Sun. Will pre-training ever end? a first step toward next-generation foundation mllms via self-improving systematic cognition. *arXiv preprint arXiv:2503.12303*, 2025.
- [45] Xiaoying Zhang, Hao Sun, Yipeng Zhang, Kaituo Feng, Chaochao Lu, Chao Yang, and Helen Meng. Critique-grpo: Advancing llm reasoning with natural language and numerical feedback. *arXiv preprint arXiv:2506.03106*, 2025.
- [46] Yipeng Zhang, Yifan Liu, Zonghao Guo, Yidan Zhang, Xuesong Yang, Chi Chen, Jun Song, Bo Zheng, Yuan Yao, Zhiyuan Liu, et al. Llava-uhd v2: an mllm integrating high-resolution feature pyramid via hierarchical window transformer. *arXiv* preprint arXiv:2412.13871, 2024.

- [47] Zhuosheng Zhang, Aston Zhang, Mu Li, Hai Zhao, George Karypis, and Alex Smola. Multimodal chain-of-thought reasoning in language models. *arXiv preprint arXiv:2302.00923*, 2023.
- [48] Yilun Zhao, Lujing Xie, Haowei Zhang, Guo Gan, Yitao Long, Zhiyuan Hu, Tongyan Hu, Weiyuan Chen, Chuhan Li, Junyang Song, et al. Mmvu: Measuring expert-level multi-discipline video understanding. *arXiv preprint arXiv:2501.12380*, 2025.
- [49] Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc Le, et al. Least-to-most prompting enables complex reasoning in large language models. *arXiv* preprint arXiv:2205.10625, 2022.

NeurIPS Paper Checklist

The checklist is designed to encourage best practices for responsible machine learning research, addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove the checklist: **The papers not including the checklist will be desk rejected.** The checklist should follow the references and follow the (optional) supplemental material. The checklist does NOT count towards the page limit.

Please read the checklist guidelines carefully for information on how to answer these questions. For each question in the checklist:

- You should answer [Yes], [No], or [NA].
- [NA] means either that the question is Not Applicable for that particular paper or the relevant information is Not Available.
- Please provide a short (1–2 sentence) justification right after your answer (even for NA).

The checklist answers are an integral part of your paper submission. They are visible to the reviewers, area chairs, senior area chairs, and ethics reviewers. You will be asked to also include it (after eventual revisions) with the final version of your paper, and its final version will be published with the paper.

The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation. While "[Yes]" is generally preferable to "[No]", it is perfectly acceptable to answer "[No]" provided a proper justification is given (e.g., "error bars are not reported because it would be too computationally expensive" or "we were unable to find the license for the dataset we used"). In general, answering "[No]" or "[NA]" is not grounds for rejection. While the questions are phrased in a binary way, we acknowledge that the true answer is often more nuanced, so please just use your best judgment and write a justification to elaborate. All supporting evidence can appear either in the main paper or the supplemental material, provided in appendix. If you answer [Yes] to a question, in the justification please point to the section(s) where related material for the question can be found.

IMPORTANT, please:

- Delete this instruction block, but keep the section heading "NeurIPS Paper Checklist",
- Keep the checklist subsection headings, questions/answers and guidelines below.
- Do not modify the questions and only use the provided macros for your answers.

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]
Justification: [NA]
Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: See Appendix E

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]
Justification: [NA]

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and crossreferenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: See the Methods section, the Experiments section and Appendix D for details Guidelines:

• The answer NA means that the paper does not include experiments.

- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
- (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
- (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: All code, models, and data will be released

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).

• Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: See the Methods section, the Experiments section and Appendix D for details Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental
 material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: The training of MLLMs is computationally expensive

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error
 of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: See the Experiments section and Appendix D for details

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.

- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

Answer: [Yes]
Justification: [NA]
Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: There are no specific negative societal impacts associated with this work Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]
Justification: [NA]
Guidelines:

• The answer NA means that the paper poses no such risks.

- Released models that have a high risk for misuse or dual-use should be released with
 necessary safeguards to allow for controlled use of the model, for example by requiring
 that users adhere to usage guidelines or restrictions to access the model or implementing
 safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do
 not require this, but we encourage authors to take this into account and make a best
 faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]
Justification: [NA]
Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: All code, models, and data will be released

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]
Justification: [NA]

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]
Justification: [NA]

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent)
 may be required for any human subjects research. If you obtained IRB approval, you
 should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]
Justification: [NA]
Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (https://neurips.cc/Conferences/2025/LLM) for what should or should not be described.

A Additional Experiments

A.1 Scaling Up RL Training

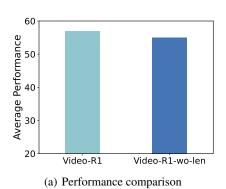
In previous experiments, our Video-R1 is trained with 1k RL steps, similar to prior works that apply the R1 paradigm in other domains [7, 29]. To further explore the impact of scaling up reinforcement learning, we extend the training to 10k RL steps. As shown in Table 3, the performance of Video-R1 generally improves on various benchmarks with more RL training steps. These results demonstrate the effectiveness of large-scale RL training and suggest that additional training can further enhance the model's reasoning capabilities.

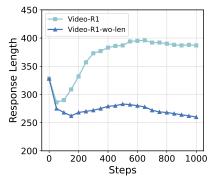
Table 3:	Paculte	on more	DΙ	training	ctone
Table 5:	Kesuits	on more	KL	training	steps.

Models	Frames	Video Reasoning Benchmark			Video General Benchmark		
		VSI-Bench	VideoMMMU	MMVU (mc)	MVBench	TempCompass	VideoMME (wo sub)
Video-R1-7B (1k steps)	16	34.6	49.8	64.2	62.7	72.6	57.4
Video-R1-7B (1k steps)	32	35.8	52.3	63.8	63.9	73.2	59.3
Video-R1-7B (1k steps)	64	37.1	52.4	63.8	64.8	73.2	61.4
Video-R1-7B (10k steps)	16	34.5	50.7	64.2	63.7	73.7	57.2
Video-R1-7B (10k steps)	32	35.7	52.7	66.2	64.7	74.2	59.3
Video-R1-7B (10k steps)	64	37.8	51.4	65.0	65.5	74.2	61.8

A.2 Effect of Length Reward Analysis

To evaluate the impact of the proposed length reward, we conduct an ablation study by comparing Video-R1 with a variant trained without the length reward, denoted as Video-R1-wo-len. As shown in Figure 8(b), removing the length reward leads to a clear decline in response length during RL training, whereas Video-R1 maintains a steadily increasing length that eventually stabilizes at a higher level. Finally, Video-R1-wo-len achieves lower average performance across benchmarks compared to Video-R1, as shown in Figure 8(a). We guess that this is because the model fails to preserve a deep reasoning style without the length reward during training, instead favoring shorter and potentially less informative responses. This likely results in weaker generalization during evaluation. These findings highlight the importance of encouraging the model to maintain a moderate level of reasoning effort during training.





(b) Length curve comparison

Figure 8: Effect of length reward.

B Reasoning Examples

Video-R1-7B: https://doi.org/10.10/

Video-R1-7B: ht

Figure 9: An example of Video-R1-7B's reasoning output on MMVU Benchmark.



Figure 10: An example of Video-R1-7B's reasoning output on MVBench.

Prompt Template

C.1 Prompt Template for Training and Inference

Figure 11 illustrates the prompt template for training and inference of all models. We also use this prompt for the COT annotation.

Prompt Template for Training and Inference

{Question}

Please think about this question as if you were a human pondering deeply. Engage in an internal dialogue using expressions such as 'let me think', 'wait', 'Hmm', 'oh, I see', 'let's break it down', etc, or other natural language thought expressions. It's encouraged to include selfreflection or verification in the reasoning process. Provide your detailed reasoning between the <think> and </think> tags, and then give your final answer between the <answer> and </answer>

{OutputTemplate}

OutputTemplate:

"multiple choice": Please provide only the single option letter (e.g., A, B, C, D, etc.) within the <answer> </answer>

"numerical": Please provide the numerical value (e.g., 42 or 3.14) within the <answer> </answer> tags.

"OCR": Please transcribe text from the image/video clearly and provide your text answer within the «answer»

"free-form": Please provide your text answer within the <answer> </answer> tags.
"regression": " Please provide the numerical value (e.g., 42 or 3.14) within the <answer> </answer> tags.

Figure 11: Prompt template for training and inference.

C.2 Prompt Template for Temporal Reasoning Evaluation

Figure 12 illustrates the prompt template for temporal reasoning evaluation.

Prompt Template for Temporal Reasoning Evaluation

{Question}

You are given a question about a video and a response to that question. Your task is to assess whether the response engages with any aspect of the video's temporal structure — that is, any elements related to time, sequence, or progression. This may include references to actions occurring in a particular order, events unfolding over time, descriptions of what happens before or after, mentions of duration, speed, or direction, or any general sense of temporal flow conveyed in the response. Consider a broad range of interpretations, including responses that reflect an understanding of how events evolve or follow one another. Provide your reasoning between the <think> and </think> tags, and then write either 'yes' or 'no' between the <answer> and </answer> tags.

Figure 12: Prompt template for temporal reasoning evaluation.

Additional Implementation Details

We use the Adam optimizer with a learning rate of 1e-6 to train our model. The SFT stage takes approximately 40 hours per epoch, while the RL stage takes around 15 hours for 1k steps. The hyperparameter β in the KL divergence term of the GRPO algorithm is set to 0.04. To ensure training stability, we apply a weight decay rate of 0.01 and clip the maximum gradient norm to 5. The maximum response length is set to 768 tokens.

E Limitations and Future Works

We envision this work as a foundation for advancing research in video reasoning with MLLMs. Below, we outline its limitations and potential avenues for future work:

- **Increasing Frames Number.** Currently, our model is trained with 16 video frames, which may limit its ability to handle long-range temporal dependencies. In future work, we can develop more efficient training and inference strategies that allow scaling to longer videos, enabling more comprehensive temporal reasoning.
- Better Temporal Modeling Method. Although T-GRPO introduces effective temporal-aware reasoning, it brings additional computational overhead due to contrastive evaluation and reward calculation. This could be mitigated through inference acceleration framework such as vLLM, or by exploring more efficient mechanisms for temporal modeling.
- Adaptive Response Length Control. Our current length control mechanism applies a fixed
 reward within a predefined range, regardless of the complexity of each sample. Future work
 could explore dynamic length control strategies, where the model adaptively determines the
 appropriate response length based on the difficulty or type of the question.
- Refined Image-to-Video Knowledge Transfer. At present, we incorporate image-based reasoning data in a straightforward manner by mixing it into the training set. Future research could design more principled approaches for leveraging image data to more effectively transfer reasoning ability from images to videos.
- Generalist Video Reward Modeling. Currently, we use rule-based reward functions tailored to different tasks. A promising future direction is to develop a generalist video reward model, which is capable of providing consistent and scalable reward signals across various video reasoning tasks. Such a model could reduce reliance on handcrafted rules and enable broader applicability of reinforcement learning in video understanding.